

Geometric VLAD for Large Scale Image Search

Zixuan Wang^{*}, Wei Di[†], Anurag Bhardwaj[†], Vignesh Jagadeesh[†], Robinson Piramuthu[†]

^{*} Dept. of Electrical Engineering, Stanford University, CA 94305

[†] eBay Research Labs, eBay Inc., San Jose, CA 95125

zxwang@stanford.edu, wedi@ebay.com, anbhardwaj@ebay.com, vjagadeesh@ebay.com, rpiramuthu@ebay.com

Abstract

We present a novel compact image descriptor for large scale image search. Our proposed descriptor - Geometric VLAD (gVLAD) is an extension of VLAD (Vector of Locally Aggregated Descriptors) that incorporates weak geometry information into the VLAD framework. The proposed geometry cues are derived as a membership function over keypoint angles which contain evident and informative information but yet often discarded. A principled technique for learning the membership function by clustering angles is also presented. Further, to address the overhead of iterative codebook training over real-time datasets, a novel codebook adaptation strategy is outlined. Finally, we demonstrate the efficacy of proposed gVLAD based retrieval framework where we achieve more than 15% improvement in mAP over existing benchmarks.

1. Introduction

Proliferation of large-scale image collections on web has made the task of efficient image retrieval challenging. Given a query image or region, the goal is to retrieve images of the same object or scene from a large scale database with high accuracy, efficiency and less memory usage. One of the core problems is how to concisely represent the visual information present in images. A number of methods have been proposed recently that address this issue from both computational efficiency as well as retrieval accuracy perspectives. However, there is a growing need for algorithms that can achieve reasonable trade-offs on both these aspects. Vector of Locally Aggregated Descriptors (VLAD) [8] proposed by Jégou et al. is one of the seminal contributions in this area as they show that compact and accurate VLAD representation is able to scale to billions of descriptors (by avoiding expensive hard disk operations) and still retain superior retrieval performance. However, one of the limitations of this representation is its inability to incorporate more keypoint level information that can potentially lead to enhanced performance. One such information is the domi-




Figure 1. gVLAD motivation: A set of key points (denoted in red dot) locates in the feature space with same distance r towards the centroid c_i , assuming they are of same feature descriptor. VLAD is unable to differentiate between the configurations shown in the left and right figures which differ only in the orientations of key-points (depicted by arrow). However, by separating keypoints into two bins according to their dominant orientation, and measuring distance of points from each bin towards the centroid separately, the proposed gVLAD can successfully differentiate between the two configurations.

nant angle of the detected keypoint, also referred to as “Keypoint Angle”, which is often discarded for the sake of obtaining rotational invariance in matches. A toy example is illustrated in Figure 1, in which VLAD is unable to differentiate between the configurations shown in the left and right figures where keypoints (red dots) differ only in their orientations, while having same descriptor representation and distance in the feature space towards the centroid c_i . Thus, we hypothesize that keypoint angles provide useful geometric cues which can be very useful for matching images. Integrating this information in a principled way can substantially improve the performance of existing VLAD based representation. In this paper, we present Geometric VLAD (gVLAD) that strengthens the VLAD representation by incorporating weak geometric cue in form of keypoint angles.

Our contributions in this paper are as follows:

- **Angle Binning Based VLAD:** We propose a novel formulation of gVLAD that incorporates low level

keypoint angles in form of a membership function into the VLAD representation.

- **Circular Preserved Angle Membership Learning:** We propose a simple but effective principled technique to learn the membership function of keypoint angles based on trigonometric transform and clustering in a fashion that preserves their circular distribution.
- **Codebook Adaptation:** To eliminate the need of iterative codebook training for large scale real-world image collections, a codebook adaptation scheme is presented.
- **Z-Score Normalization:** Z-score based normalization technique is proposed that outperforms existing normalization methods for VLAD-based representations.
- **Superior New Benchmark Results:** State-of-the-art image retrieval performance of proposed framework over a number of existing retrieval benchmarks are achieved.

The paper is organized as follows. In section 2, we outline related work in large-scale image search and strategies of integrating geometric information into image representations. In section 3, we describe the geometric VLAD representation in detail. In section 4, we demonstrate the performance gain on *Oxford*, *Holidays* and *Paris* benchmarks, as well as on extended large scale datasets. We conclude the paper and discuss the future work in section 5.

2. Related Work

The Bag-of-Words (BoW) representation is one of the most widely used method for image retrieval [16, 14]. It quantizes each local descriptor SIFT [11] or SURF [3], to its nearest cluster center and encodes each image as a histogram over cluster centers also known as “Visual Words”. Good retrieval performance is achieved with a high dimensional sparse BOW vector, in which case inverted lists can be used to implement efficient search. However, the search time grows quadratically as the number of images increase [4].

To overcome this issue, the Fisher kernel based approach proposed by Perronnin *et al.* [13] transforms an set of variable-sized independent samples into a fixed size vector representation. The samples are distributed according to a parametric generative model, in this case a Gaussian Mixture Model (GMM) estimated on a training set. A simplified version of Fishers kernels, the VLAD is proposed by Jégou *et al.* [8, 9]. It encodes the difference from the cluster center in a more direct manner, rather than the frequency assigned to the cluster. It requires less computation than Fisher kernels but can achieve comparable retrieval performance.

However, most of existing methods ignore the geometric information present in images. Spatial re-ranking [14]

is usually used as a geometric filter to remove unrelated images from retrieval results. However, due to its expensive computation it is applied only to top ranked images for re-ranking. The spatial pyramid [10] is a simple extension of the BOW representation which partitions the image into increasingly fine sub-regions and computes histograms of local features found inside each sub-region. It shows improved performance on scene classification tasks. The weak geometric consistency constraints (WGC) [7] uses angle and scale information from key points to verify the consistency of matching descriptors. It can improve the retrieval performance significantly. Recently, Zhang *et al.* [17] propose a technique to encode more spatial information through the geometry-preserving visual phrases (GVP) which requires a pair of images to obtain geometric information. Chum *et al.* [5] propose geometric min-hash, which extends min-hash by adding local spatial extent to increase the discriminability of the descriptor. It can be used for nearly duplicate image detection but has not achieved the state-of-the art performance in retrieval.



Figure 2. Each figure on the left shows an input image. Each figure on the right shows the detected key points in the image. Keypoints are grouped into four bins based on their angles represented by the direction of the line, and is colored by four unique colors. Length of the line corresponds to scale. Note that each image has a distinct representation based on the orientation of key points which suggests that this information can be potentially useful in image representation.

3. Proposed Framework

In this section, we introduce the Geometric VLAD (gVLAD) to improve retrieval performance by incorporating low level angle information from the key points into VLAD framework.

3.1. Geometric VLAD

Let us represent the local descriptor \mathbf{x} to be d -dimensional vector (e.g. SURF or SIFT descriptors). Codebook or visual words are denoted as $\mu = [\mu_1, \mu_2, \dots, \mu_K]$, where K represents the size of the vocabulary. Let $NN(\mathbf{x})$ represent the nearest-neighbor function that maps an input descriptor \mathbf{x} to its nearest visual word i where $1 \leq i \leq K$. In the original VLAD [8, 9], to represent a given image, a set of local descriptors are extracted first. Then, the contribution of each visual word v_i is defined by accumulating distances of all the descriptors that belong to the i^{th} visual word μ_i as:

$$v_i = \sum_{\mathbf{x}: NN(\mathbf{x})=i} \mathbf{x} - \mu_i \quad (1)$$

Such representation is further L2-normalized, and concatenated to form a vector representation with size $d \times K$ to represent each image. However, the above formulation suffers from the drawback that it is unable to incorporate extra descriptor level information such as angle which can be of very useful in providing a weak geometrical cue. Thus, we present a gVLAD representation which encodes such angle information of the descriptor into the VLAD framework for efficient image matching. To define gVLAD, we redefine a local descriptor by \mathbf{x}^θ , where \mathbf{x} still represents the appearance feature vector of the descriptor and θ represents the angle of the descriptor, i.e. the dominant angle of the keypoint. For example, in SIFT descriptor, the angle corresponds to the dominant direction of gradient within a local window. To model the distribution of angles, we introduce clustering idea and define a membership function over the angles as: $\psi(\theta(\mathbf{x})) : 0 \leq \theta < 2\pi \rightarrow \{1, 2, \dots, M\}$, where M denotes the total number of angular bins.

The gVLAD v_i^j for i^{th} of the K visual words (feature bin) and j^{th} of the M angular bins can now be represented as:

$$v_i^j = \begin{cases} \sum_{\mathbf{x}^\theta: NN(\mathbf{x})=i} \mathbf{x}^\theta - \mu_i & \text{if } \psi(\theta(\mathbf{x})) = j \\ \mathbf{0}^d & \text{if } \psi(\theta(\mathbf{x})) \neq j \end{cases} \quad (2)$$

where d is the dimension of feature vector of local descriptor \mathbf{x} . The contribution of each visual word V_i in the geometric VLAD can now be written as combining individual contributions from each angle bin:

$$V_i = [v_i^1, v_i^2, \dots, v_i^{M-1}, v_i^M] \quad (3)$$

where V_i is a row vector with size of $d \times M$. Our geometric VLAD (gVLAD) representation \mathcal{V} is defined by accumulating contributions of from all K visual words, and has D dimensions: $D = K \times d \times M$.

$$\mathcal{V} = [V_1, V_2, \dots, V_{K-1}, V_K] \quad (4)$$

3.2. Learning Membership Function - $\psi(\theta)$

One principal way to learn the membership function $\psi(\theta)$ is to apply clustering over the angle distribution and find the appropriate membership assignments for each angle value among M learned clusters. Typically angles have a circular distribution of in the range of $[0, 2\pi)$, whereas existing clustering algorithms that based on L2 distance such as k -means assume a Cartesian co-ordinate space for input data, and can not be applied directly. To address this issue, we propose to represent each keypoint as (r, θ) , where r is the radial coordinate. Since we are only interested in angles of key point θ , we fix r as an arbitrary number $r > 0$. We now perform a non-linear transform from this polar co-ordinate to 2D Cartesian co-ordinate space using the trigonometric functions:

$$x = r \times \cos \theta \quad (5)$$

$$y = r \times \sin \theta \quad (6)$$

Thus, each angle is mapped to a point $\mathbf{z}(\theta) = (x, y)$ in this 2-d space. To learn the membership of function $\psi(\theta)$, we perform k -means clustering in this space satisfying:

$$\arg \min_{\{\alpha_1, \dots, \alpha_M\}} \sum_{i=1}^M \sum_{\mathbf{z}_j \in \Xi_i} \|\mathbf{z}_j - \alpha_i\|^2 \quad (7)$$

where α_i is the cluster centroid by averaging all points in cluster set Ξ_i . The membership of each angle θ can be estimated through:

$$\psi(\theta) = \arg \min_{i \in \{1, 2, \dots, M\}} \|\mathbf{z}(\theta) - \alpha_i\|^2 \quad (8)$$

3.3. Codebook Adaptation

Most real-world image databases grow continuously which leads to frequent codebook training processes that are often desirable. We propose a simple codebook adaptation process that can adapt existing codebooks with incremental dataset and alleviate the need of frequent large-scale codebook training. Secondly, this technique also allows codebook training from diverse datasets as codebook trained on one dataset (i.e. Paris building images) can be adapted to retrieve images from another dataset (i.e. Flickr holidays images). To define our codebook adaptation, let us represent a source dataset of images S where an initial codebook $\mu = [\mu_1, \mu_2, \dots, \mu_K]$ is trained. Given a new domain specific dataset T , our goal is to adapt μ to another domain

specific codebook $\hat{\mu}$ given as:

$$\hat{\mu}_i = \frac{1}{N} \sum_{t=1}^N \gamma_i(t), \mathbf{x}^\theta(t) \in T \quad (9)$$

$$\text{where } \gamma_i(t) = \begin{cases} \mathbf{x}^\theta(t) & \text{if } NN(\mathbf{x}^\theta(t)) = \mu_i \\ \mathbf{0}^d & \text{if } NN(\mathbf{x}^\theta(t)) \neq \mu_i \end{cases} \quad (10)$$

where N is the total number of descriptors in dataset T and $\mathbf{x}^\theta(t)$ represents t^{th} descriptor. In our experiment, the initial codebook μ is trained using the *Paris* dataset. For all the other experiments on different datasets, $\hat{\mu}_i$ is used in conjunction with Equation 2 to compute the representation of the geometric VLAD.

3.4. gVLAD Normalization

Normalization is important to effectively and correctly measure distance between vector representation. Here we propose three stages of normalization. First, we use the intra-normalization [2], where the sum of residuals of each visual word v_i^j is L2 normalized independently, where $1 \leq i \leq K$ and $1 \leq j \leq M$. This step is followed by inter-Z-score normalization across different visual words. Given a vector X , its Z-score normalization is computed as: $\frac{X - \mu}{\sigma}$, where μ and σ represent the mean and standard deviation of X . Let's denote the t^{th} entry of V_i as $v_{i,t}$, where V_i is defined in Equation 3. We apply the inter-Z-score normalization on each $[v_{1,t}, v_{2,t}, \dots, v_{i,t}, \dots, v_{K,t}]$, where $1 \leq t \leq M \times d$ and $1 \leq i \leq K$. At last, the gVLAD vector \mathcal{V} is L2 normalized $\mathcal{V} := \mathcal{V} / \|\mathcal{V}\|_2$.

3.5. PCA Whitening

Given a large collection of images, the size of representation needs to be carefully considered so as to be feasible for practical real time retrieval. For instance, using only 256 visual words with 64 dimensional SURF descriptors and 4 angle bins generates a feature representation of size $D = 64 \times 256 \times 4 = 65,536$. To achieve memory-efficient representation of this vector, we use standard PCA with pre-whitening as described in [6]. The PCA whitening matrix can be expressed in the form of:

$$\mathbf{P} = \mathbf{D}^{-1/2} \mathbf{E}^T \quad (11)$$

where $\mathbf{E} \mathbf{D} \mathbf{E}^T = E\{\bar{\mathbf{V}} \bar{\mathbf{V}}^T\}$ is the eigenvector decomposition of the covariance matrix of the (zero mean) data $\bar{\mathbf{V}}$, where each row $\bar{\mathbf{V}}_l = \mathcal{V}_l - \mathcal{V}_0$, and \mathcal{V}_0 is the mean vector computed from all gVLAD representation vectors. $\mathbf{D} = \text{diag}[d_1, d_2, \dots, d_D]$ is the $D \times D$ diagonal matrix containing the eigenvalues and $\mathbf{E} = [e_1, e_2, \dots, e_D]$ is an orthogonal $N \times D$ matrix having the eigenvectors as columns.

The obtained whitened gVLAD representation is:

$$\tilde{\mathcal{V}}_l = \mathbf{P}(:, 1 : \rho)^T \times \bar{\mathcal{V}}_l \quad (12)$$

where ρ is the number of eigenvectors to keep, i.e. the dimension of reduced feature vectors. $\tilde{\mathcal{V}}_l$ is then L2 normalized. The complete algorithm is outlined in Algorithm 1.

Algorithm 1 Computation of gVLAD descriptor

- 1: **S1:** Keypoint detection and descriptor: compute image descriptors \mathbf{x}^θ , where \mathbf{x} is the appearance vector, and θ represents the angle.
 - 2: **S2:** Generate visual vocabulary $[\mu_1, \mu_2, \dots, \mu_K]$ using k -means on all descriptors from training data.
 - 3: **S3:** Learning membership function $\psi(\theta)$ for each \mathbf{x}^θ
 - 4: $\arg \min_{i \in \{1, 2, \dots, M\}} \|\mathbf{z}(\theta) - \alpha_i\|^2$
 - 5: **S4:** Compute geometric VLAD v_i^j :
 - 6: $v_i^j = \begin{cases} \sum_{\mathbf{x}^\theta: NN(\mathbf{x})=i} \mathbf{x}^\theta - \mu_i & \text{if } \psi(\theta(\mathbf{x})) = j \\ \mathbf{0}^d & \text{if } \psi(\theta(\mathbf{x})) \neq j \end{cases}$
 - 7: $V_i = [v_i^1, v_i^2, \dots, v_i^{M-1}, v_i^M]$
 - 8: $\mathcal{V} = [V_1, V_2, \dots, V_{K-1}, V_K]$
 - 9: **S5:** Codebook Adaptation:
 - 10: $\hat{\mu}_i = \frac{1}{N} \sum_{t=1}^N \gamma_i(t), \mathbf{x}^\theta(t) \in T$
 - 11: **S6:** Intra-normalization, Inter-Z-score normalization, and L2-normalization.
 - 12: **S7:** PCA whitening:
 - 13: $\tilde{\mathcal{V}}_l = \mathbf{P}(:, 1 : \rho)^T \times \bar{\mathcal{V}}_l$
-

4. Experiments & Evaluations

4.1. Benchmark Dataset

We evaluate the proposed approach on several public available benchmark datasets: *Oxford* buildings, *Paris* and *Holidays*. Large scale experiments are conducted on these datasets by adding 1M Flickr images as distractors [7]. For each of these datasets, performance is measured by mean average precision (mAP) over a set of pre-defined queries and their annotated ground truth matches.

Holidays Dataset: *Holidays* dataset [7] contains 1491 high resolution personal holiday photos with 500 annotated queries. For large scale experiments, 1 million Flickr images are added to it to create *Holidays + Flickr 1M* dataset. About 5%-10% of the images in holiday dataset have orientations which are unnatural for human observer [12]. We manually rotate these images to create *Rotated Holidays* dataset.

Oxford Dataset: This dataset, *Oxford 5K* contains 5062 images of Oxford buildings gathered from Flickr [14]. There are 55 query images each with a rectangular bounding box specifying the region of interest. To test large scale retrieval, it is firstly extended with a 100K Flickr images¹ to create *Oxford 105K* dataset. We further extend the dataset with 1 million Flickr image² creating *Oxford 5K + Flickr*

¹<http://www.robots.ox.ac.uk/vgg/data/oxbuildings/flickr100k.html>

²<http://press.liacs.nl/mirflickr/>



Figure 3. The pipeline of the geometric VLAD (gVLAD) descriptor computation.

1M dataset.

Paris Dataset: The Paris Dataset [15] *Paris 6K* consists of 6412 images collected from Flickr by searching for particular Paris landmarks. There are 60 query images, each with a rectangular bounding box specifying the region of interest. We found that both 100K Flickr images and Flickr 1M images contains a large number of Paris landmarks, hence we do not extend the Paris dataset with Flickr images.

4.2. Implementation Details

Descriptor computation: The pipeline of computing gVLAD descriptor is characterized in Figure 3. First, all images are resized to 1024×768 . We find that when using the original resolution of *Holidays* images, the performance is inferior to the down-sampled images. We can also benefit from the speed when using smaller images. In *Oxford* and *Paris* datasets, bounding boxes are provided for queries. We only extract descriptors inside bounding boxes. We use the SIFT and SURF implementations in OpenCV³ to detect keypoints and extract descriptors. Each SIFT descriptor has 128 dimensions and each SURF descriptor has 64 dimensions. We find that VLAD based features have better performance using SURF keypoints and descriptors [3] than SIFT keypoints and descriptors [11]. In general, we observed about 10% improvement using SURF as compared to SIFT. More details about the performance difference can be seen by comparing results in Table 2 and Table 3.

Angle Membership Function: The angle distribution of SURF keypoints from *Holidays* dataset is shown in Figure 4 (a). We find that majority keypoints have vertical or horizontal angles as detectors have larger response at these points, resulting in roughly 4 centers ($\pi/2, \pi, 2\pi/3, 2\pi$). To learn the membership function of each keypoint angle, we apply the proposed approach in 3.2. Because larger number of bins will increase the dimension of the final derived gVLAD feature, to gain a reasonable representation as well as low dimensionality, we set the number of angle bins to be 4 to fit the distribution. A $\pi/4$ offset and a set of evenly distributed bins: $[-\pi/4, \pi/4)$, $[\pi/4, 3\pi/4)$, $[3\pi/4, 5\pi/4)$ and $[5\pi/4, 7\pi/4)$, are automatically estimated from the algorithm, which is visualized in Figure 4 (b). We use this angle bin partition in following evaluations. We had also experimented using different number of bins and offset on *Rotated Holidays* dataset. We observe increasing performance as more bins are used, as shown in Figure 5. This

is because that increasing number of bins is equivalent to increasing number of subspaces, by which the distance of descriptor towards centroid can be computed in a more discernible way. However, gains by using 5 or 6 bins as compared to the predicted angle partition (4 bins with $\pi/4$ offset) by propose algorithm are marginal, also our learned setting has much smaller dimensions.




Figure 4. (a) Distribution of keypoint angles from *Holidays* dataset. (b) Learnt 4 angle bins with $\pi/4$ offset.




Figure 5. Performance using different angle bins and offset on *Rotated Holidays* dataset. Based on empirical observation, we apply angle membership learning with $M = 4$, and use the learned partition in our experiment. Such setting gives near best performance but lower dimension of gVLAD features as compared to using more bins for modeling angle distribution.

Vocabulary Generation: The vocabulary consisting of $K = 256$ visual words is computed from all SURF descriptors on Paris dataset. Various different cluster initializations of k -means are executed and the best clustering is used as the vocabulary for all evaluations. As the number

³<http://opencv.org/>

	Angle Bins					
Q	2	4	8	18	36	72
mAP	0.015	0.037	0.149	0.241	0.261	0.269

Table 1. Retrieval performance on *Rotated Holidays* dataset using only Angle binning histogram with varying dimensions, no appearance information is used. Given only 72 dimension of angle histogram, a surprising 26.9% mAP result is achieved.

of extracted descriptors is typically much larger than K , e.g. even the smallest Holidays dataset contains 8.3 million SURF descriptors. This vocabulary can be considered independent from all datasets. Such simplification has been used in literature [2]. For every dataset, this vocabulary is used as a reference vocabulary and a vocabulary adaptation is performed as described in section 3.3.

Retrieval: During retrieval, L2 distance is computed to rank images with respect to input query. Since our focus is generating a compact and efficient image descriptor, to illustrate the power of the proposed descriptor, we use brute-force distance computation to report our results. However, our proposed descriptors can in principle be used with approximate distance matching or other hashing based techniques as well, which is beyond the scope of this paper.

4.3. Performance Evaluation & Analysis

The performance in all retrieval experiments is evaluated using the mean average precision (mAP), which is defined as the mean of the average precision over all the queries given a dataset. Average Precision is computed as the average of the precision value obtained for the set of top k images after each relevant image is retrieved. We use standard evaluation packages obtained from the data websites.

The Power of Angle: To illustrate the power of angle, we performed a simple experiment which uses only the angle information from each keypoint to retrieve similar images. After obtaining the keypoints and the angle of each keypoint, we generate an angle histogram for each of image by binning all angles into Q bins. We use L2 distance to compute the similarity between angle descriptors. Table 1 shows the retrieval results on *Rotated Holidays* dataset. It can be seen that surprisingly using only angle information (without any appearance information from SURF or other descriptors), we can still achieve about 26.9% mAP results. Note that the dimension of the angle bin histogram for the best result is only 72, which is a much smaller number compared to conventional BOW or VLAD descriptors.

Step-by-Step Performance Evaluation: To show the performance gain obtained from each of the proposed steps, we performed a step-by-step experiment on *Rotated Holidays* dataset and baselined it with VLAD performance. Results are listed in Table 3. All results use SURF detector and SURF descriptor. It can be seen that by adding inter-Z-Score normalization to the original VLAD, the performance

Dataset	VLAD*	gVLAD
Holidays	0.548	0.710
Rotated Holidays	0.550	0.786

Table 2. Comparison of our proposed gVLAD with VLAD on benchmark datasets. Both VLAD and gVLAD use SIFT detectors and root SIFT descriptors. 256 visual words are used. The feature dimension of VLAD is $256 \times 128 = 32768$, and for gVLAD is $256 \times 128 \times 4 = 131072$. * denotes our implementation.

is increased by 5.4%. Performing Angle Binning over VLAD leads to a gain of 7.3%. By combining both Angle Binning and Z-Score normalization, we achieve 14.7% improvement over VLAD representation. Performing vocabulary adaptation for *Rotated Holidays* dataset provides additional 3.8% performance gain. Finally, PCA whitening is applied which is able to reduce the dimension significantly with only about 1.1% performance loss, as compared to PCA without whitening having a loss of 3.6%. To demonstrate the performance of low-dimensional gVLAD descriptor using PCA whitening, we further plot the mAP performance curve by varying $\rho = 2^4$ to 2^{16} in Figure 6. It can be seen that with only 32 dimensions, the performance by the proposed descriptor can reach to $mAP = 0.737$, which already outperforms the original VLAD descriptor using 1024 visual words with 65,536 dimension ($mAP = 0.670$) as shown in Table 3.

We also test our proposed method using SIFT detectors and root SIFT [1] descriptors, since most previous published work use SIFT. For fair comparison, we implement VLAD with root SIFT descriptors, which have better performance compared with 0.526 on *Holidays* dataset reported in [8]. Results as shown in Table 2 demonstrate the superior performance of proposed approach over SIFT descriptors as well. As noted, comparing Tabel 2 and Table 3, we observe in general that using SURF descriptors outperforms SIFT based descriptors.




Figure 6. Dimension reduction on original gVLAD descriptor using PCA whitening. The original feature dimension is 65,536. After compressed to 128-D, the mAP decreases only about 1%.

Method	Dimension	mAP
VLAD ($K = 256$)* [9]	16,384	0.662
VLAD ($K = 1024$)* [9]	65,536	0.670
VLAD ($K = 256$) + inter-norm	16,384	0.716
VLAD ($K = 256$) + Angle Binning	65,536	0.735
+ inter-norm	65,536	0.809
+ Voc Adaptation	65,536	0.847
PCA	128	0.811
PCA + whitening	128	0.836

Table 3. mAP on the *Rotated Holidays* dataset comparing to start-of-art results. Best performances are in bold. *VLAD result in this table are based on our implementation. All results use SURF detector and SURF descriptor.

Method	Dimension	Holidays	Oxford
BoW 20k-D [9] [16]	20,000	0.452	0.354
BoW 200k-D [9] [16]	200,000	0.540	0.364
Improved Fisher [13]	16,384	0.626	0.418
VLAD [8]	8,192	0.526	-
VLAD + SSR [9]	16,384	0.598	0.378
Improved VLAD + SSR [2]	32,768	-	0.532
VLAD + intra-norm [2]	32,768	0.646	0.555
Ours	65,536	0.812	0.626

Table 4. mAP performance by full size gVLAD descriptors as compared to state-of-the-art results on *Holidays* and *Oxford*. Existing approaches are based on SIFT descriptors, while the proposed gVLAD descriptor uses SURF detector and SURF descriptor. Best performances are in bold.

Full Size & Low Dimensional gVLAD Descriptors: We compared our proposed method with several benchmark results in [9, 16, 13, 2] for both full size descriptor and dimension reduced descriptor ($\rho = 128$). Experiments are done using both *Holidays* dataset and *Oxford 5K* dataset. Table 4 shows that the proposed approach significantly outperforms the state-of-the-art performance by approximately 16.6% and 7% on *Holidays* and *Oxford 5K* dataset respectively. For low dimensional case, as shown in Table 5, our algorithm outperforms the best state-of-art result by 15% on both datasets.

Comparing Table 4 and 5, results also show that the proposed gVLAD descriptor is quite powerful in the sense that even with PCA whitening and reduced dimension, it can still achieve better result as compare to the best benchmark results with full size descriptors. In addition, PCA whitening based dimension reduction only results in small amount of performance decrease which is about 2.76% in average of both datasets, and 1.1% in the best case (*Rotated Holidays*).

Performance on Large Scale Dataset: We scale the proposed algorithm to large scale image dataset with millions of images, and test on both using full size gVLAD and PCA dimension reduced 128-D descriptors. In total, 4 large scale datasets are used, including *Holidays + Flickr 1M*, *Rotated Holidays + Flickr 1M*, *Oxford 105K*, and *Oxford*

Method	Holidays	Oxford
GIST [9]	0.365	-
BoW [9, 16]	0.452	0.194
Improved Fisher [13]	0.565	0.301
VLAD [8]	0.510	-
VLAD + SSR [9]	0.557	0.287
Multivoc-BoW [6]	0.567	0.413
Multivoc-VLAD [6]	0.614	-
VLAD + intra-norm [2]	0.625	0.448
Ours	0.779	0.600

Table 5. mAP performance by gVLAD low dimensional descriptors ($\rho = 128$): comparison with state-of-the-art on the *Holidays* and *Oxford 5k* benchmarks. The existing approaches are based on SIFT descriptors, while the proposed gVLAD descriptor uses SURF detector and descriptor. Best performances are in bold.

5K + Flickr 1M. As can be seen from Table 6, our methods outperform all current state-of-the-art methods. For example, using dimension reduced 128-D gVLAD descriptors, on *Holidays + Flickr 1M* dataset, our algorithm outperforms the best result [2] reported in literature with a significant gain of 22.8%. On *Oxford 105K* dataset, we are able to achieve 11.6% better result than [2].

Further, same with our previous observation in Table 5 as compared to Table 4, Table 6 also shows performance only drops very slightly using the proposed PCA whitening. This implies that the proposed gVLAD descriptor is quite powerful. Also, being combined with proper dimension reduction schema, effective representation with computational efficiency can be achieved.

4.4. Time Complexity and Memory Footprint

Each image takes 512 bytes in memory after being converted to 128 dimensional gVLAD feature vector by PCA compression. The largest dataset (*Holidays + Flickr 1M*) in our experiment occupies 0.5GB of RAM for keeping all features in memory. To evaluate the time complexity of each step in the proposed gVLAD computation, we conduct experiments on this dataset using a Ubuntu machine with two Xeon X5675 CPUs at 3.07GHz, with 12 physical cores and 24 logical cores in total. We rely on multi-threading whenever possible. Table 7 illustrates the average results on 10 randomly selected queries. As shown, our proposed technique takes approximately 100 millisecond to compute gVLAD representation, and 750 millisecond to perform an end-to-end brute-force retrieval over the entire inventory. Since our proposed descriptors can in principle be used with other approximate distance matching or indexing schema, better retrieval speed can be expected, which will be very useful in practical applications.

Dataset	State of the Art Original Dimension	State of the Art 128-D	Ours Original Dimension	Ours 128-D	Ours Loss in PCA
<i>Holidays</i>	0.646 [2]	0.625 [2]	0.812	0.779	0.033
<i>Holidays + Flickr 1M</i>	-	0.378 [2]	-	0.607	-
<i>Rotated Holidays</i>	-	-	0.847	0.836	0.011
<i>Rotated Holidays + Flickr 1M</i>	-	-	-	0.654	-
<i>Oxford 5K</i>	0.555 [2]	0.448 [2]	0.626	0.600	0.026
<i>Oxford 105K</i>	-	0.374 [2]	-	0.490	-
<i>Oxford 5K + Flickr 1M</i>	-	-	-	0.438	-
<i>Paris 6K</i>	0.494 [15]	-	0.631	0.592	0.039

Table 6. mAP Performance on large scale datasets: comparisons with benchmark results. Best performances are in bold.

Process	Mean \pm std. (ms)
SURF detection & description	373.5 \pm 69.1
gVLAD computation	71.7 \pm 20.3
PCA compression	28.0 \pm 3.6
Nearest neighbor search	266.7 \pm 36.3

Table 7. Speed analysis based on 10 random query images from *Holidays + Flickr 1M* dataset.

5. Conclusion

We present gVLAD which is a novel extension of popular VLAD descriptor for large scale image search. Our proposed descriptor extends VLAD by integrating weak geometric cues in form of key point angles. A principled technique to represent this information as membership function over angles is also presented. The vocabulary adaptation and inter-Z-score normalization are also proposed to improve the performance of the system. Extensive experiments are conducted on existing publicly available benchmark datasets which demonstrate the superior performance of our approach. Our future work focuses on exploring efficient indexing strategies to avoid the brute-force matching of images. We are also investigating other related low level information that can be further integrated into gVLAD to make the representation more powerful.

References

- [1] R. Arandjelovic and A. Zisserman. Three things everyone should know to improve object retrieval. In *CVPR, 2012*, pages 2911–2918. IEEE, 2012. 6
- [2] R. Arandjelovic and A. Zisserman. All about vlad. In *CVPR 2013*, 2013. 4, 6, 7, 8
- [3] H. Bay, T. Tuytelaars, and L. Van Gool. Surf: Speeded up robust features. In *ECCV 2006*, pages 404–417. Springer, 2006. 2, 5
- [4] O. Chum and J. Matas. Large-scale discovery of spatially related images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(2):371–377, 2010. 2
- [5] O. Chum, M. Perdoch, and J. Matas. Geometric min-hashing: Finding a (thick) needle in a haystack. In *CVPR 2009*, pages 17–24. IEEE, 2009. 2
- [6] H. Jégou and O. Chum. Negative evidences and co-occurrences in image retrieval: The benefit of pca and whitening. In *ECCV 2012*, pages 774–787. Springer, 2012. 4, 7
- [7] H. Jegou, M. Douze, and C. Schmid. Hamming embedding and weak geometric consistency for large scale image search. In *ECCV 2008*, pages 304–317. Springer, 2008. 2, 4
- [8] H. Jégou, M. Douze, C. Schmid, and P. Pérez. Aggregating local descriptors into a compact image representation. In *CVPR 2010*, pages 3304–3311. IEEE, 2010. 1, 2, 3, 6, 7
- [9] H. Jégou, F. Perronnin, M. Douze, C. Schmid, et al. Aggregating local image descriptors into compact codes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(9):1704–1716, 2012. 2, 3, 7
- [10] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *CVPR 2006*, volume 2, pages 2169–2178. IEEE, 2006. 2
- [11] D. G. Lowe. Object recognition from local scale-invariant features. In *Computer Vision, 1999. Proceedings. Seventh IEEE International Conference on*, volume 2, pages 1150–1157. IEEE, 1999. 2, 5
- [12] M. Perdoch, O. Chum, and J. Matas. Efficient representation of local geometry for large scale object retrieval. In *CVPR 2009*, pages 9–16. IEEE, 2009. 4
- [13] F. Perronnin, Y. Liu, J. Sánchez, and H. Poirier. Large-scale image retrieval with compressed fisher vectors. In *CVPR 2010*, pages 3384–3391. IEEE, 2010. 2, 7
- [14] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *CVPR 2007*, pages 1–8. IEEE, 2007. 2, 4
- [15] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. Lost in quantization: Improving particular object retrieval in large scale image databases. In *CVPR 2008*, pages 1–8. IEEE, 2008. 5, 8
- [16] J. Sivic and A. Zisserman. Video google: A text retrieval approach to object matching in videos. In *9th IEEE International Conference on Computer Vision*, pages 1470–1477. IEEE, 2003. 2, 7
- [17] Y. Zhang, Z. Jia, and T. Chen. Image retrieval with geometry-preserving visual phrases. In *CVPR 2011*, pages 809–816. IEEE, 2011. 2